111th MPEG Geneva, Switzerland, 16 - 20 February 2015, Meeting Report
Panos Kudumakis
qMedia, Queen Mary University of London

Contents

1   Publish/Subscribe Application Format (PSAF)

PSAF (ISO/IEC 23000-16) updated in collaboration with MCO/CEL & MPEG-UD groups and reached Committee Draft status at 111th MPEG meeting.

Publish/Subscribe (PubSub) is an established communication paradigm where senders do not communicate information directly to the intended receivers but rely instead on a service that mediates the relationship between senders and receivers. In PubSub senders (called Publishers) post information to and receivers (called Subscribers) declare their interest in information to a service.

It should be noted that "Publisher" is not necessarily a user who distributes media to end users and "Subscriber" is not necessarily synonymous with end user or consumer. Indeed Publisher may represent a user who has created a media item and wants to make it available to publishers and Subscriber may represent a publisher who is looking for media to be distributed. Therefore "Publisher" could very well be a creator who announces the availability of his latest work and "Subscriber" could be a publisher who is looking for new works to publish. The mediator service is called Match Service Provider. The following walkthrough and Figure 1 describe the Publish/Subscribe mechanism.

- Publisher
  - Stores Resource
  - Creates RI
  - Stores RI
  - Creates PI
  - Publishes PI
- Subscriber
  - Creates SI
  - Subscribes SI
- Match Service Provider
  - Finds Match between Subscriber's SI and Publisher's PI
- Notifies Publisher and Subscriber
- Subscriber
  - Retrieves RI
  - Retrieves and plays Resource

Terminology: Resource Information (RI), Publication Information (PI), Subscription Information (SI) and Notification Information (NI).
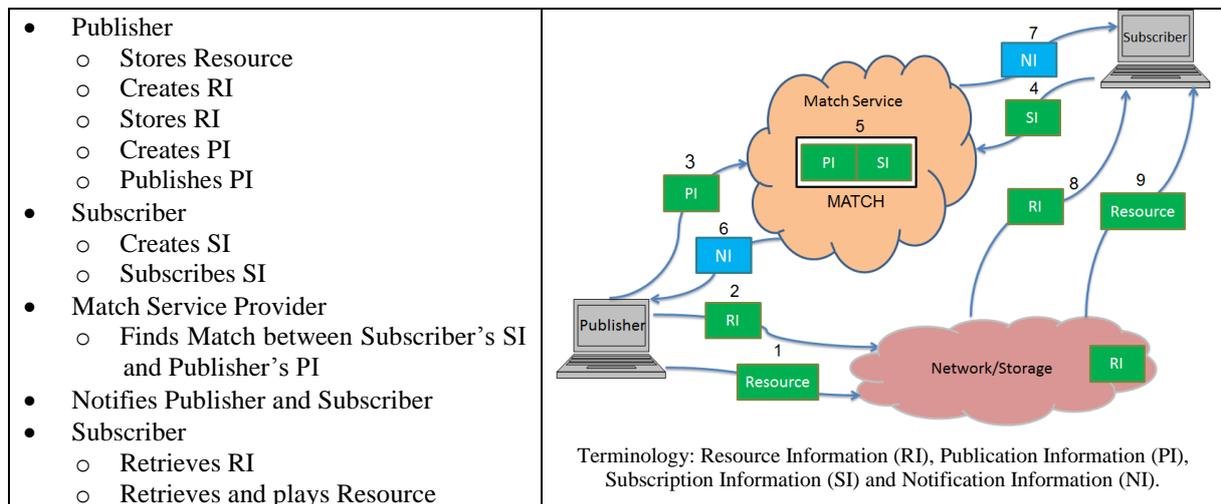
Figure 1: A Publish/Subscribe workflow for multimedia applications

The relationships between the Users of multimedia Publish/Subscribe are depicted in Figure 2.
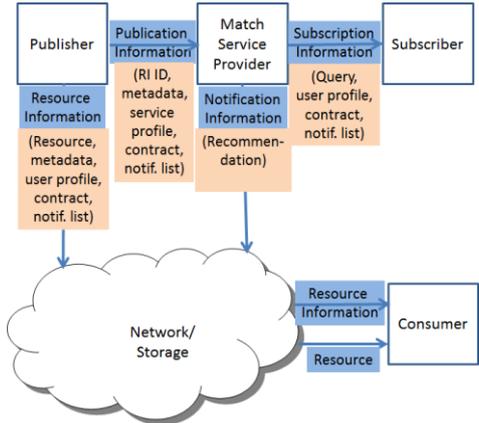

Figure 2: Relationships between users of Publish/Subscribe

Table 1 lists some of the acts performed by Publishers, Subscribers, Match Service Providers and Consumers as a result of the relationships between them.

Table 1 – Possible acts in multimedia Publish/Subscribe

Party 1	Act	Party 2	Act
Publisher	Requests match of pub with subs	MSP	E.g computes statistics
Subscriber	Requests match of sub with pubs	MSP	E.g computes statistics
Publisher	Grants right to use (e.g. play) resource	Consumer	E.g. pays for play

Requirements

1. Publishers and Subscribers shall be able to define Users or groups of Users to be/not to be notified
2. Publishers shall be able to define a standard information package (Resource Information) containing at least the following information elements
    a. Identifiers of Resource
    b. Descriptions of Resource
    c. Types of Users for which the Resource should be of interest
    d. Permissions, obligations and prohibitions for the use of the Resource by users of the Resource
    e. Request to notify a list of users that a specific use of a Resource has been made
3. Publishers shall be able to define a standard information package (Publication Information) containing at least the following information elements
    a. Metadata related to Resource Information
    b. Information related to the multimedia service provided, the target users and the recommended context of use
    c. Permissions, obligations and prohibitions for the use of Metadata related to Resource Information by the Match Service Provider
    d. Request to notify/not to notify a list of users of a Match between this Publication and Subscriptions
4. Subscribers shall be able to define a standard information package (Subscription Information) containing at least the following information elements
    a. Query related to Metadata of Resource Information
    b. Metadata related to Subscriber and the context in which the Resource sought will be consumed
    c. Permissions, obligations and prohibitions for the use of the Query by the Match Service Provider
    d. Request to notify/not to notify a list of users of a Match between this Subscription and Publications

5. Match Service Providers shall be able to define a standard information package (Notification Information) containing at least Recommendation of a Resource Information ID
6. Publishers and Subscribers shall be able to
    a. Select the Match Service Provider(s) that shall perform match of their Publications and Subscriptions
    b. Define the period of time during which the Match Service Provider shall notify Matches of their Publications or Subscriptions
    c. update their respective Publications and Subscriptions
7. Publishers, Subscribers and Match Service Providers shall be able to guarantee the authenticity of their Resource, Publication, Subscription and Notification Information

Output documents
N15115 - Requirements for Publish and Subscribe Application Format
N15201 - Request for subdivision of ISO/IEC 23000-16 Publish/Subscribe Application Format
N15202 - Text of ISO/IEC CD 23000-16 Publish/Subscribe Application Format
N15191 - Request for ISO/IEC 21000-20 2nd edition Contract Expression Language
N15192 - Text of ISO/IEC CD 21000-20 2nd edition Contract Expression Language
N15193 - Request for ISO/IEC 21000-21 2nd Media Contract Ontology
N15194 - Text of ISO/IEC CD 21000-21 2nd Media Contract Ontology
N15195 - Request for subdivision of ISO/IEC 21000-22 User Description
N15196 - Text of ISO/IEC CD 21000-22 User Description
N15197 - WD of Implementation Guidelines of User Description

2   Media Linking Application Format (MLAF)

The notion of *bridget*
Companion screen applications let user enjoying broadcast programmes access related information on other – typically internet-connected – devices. A *bridget* is a link from the programme being watched to external interactive media elements such as web pages, images, audio clips, different types of video (2D, multi-view, with depth information, free viewpoint) and synthetic 3D models.
A *bridget* can be just a link from a portion of a "source" programme to a single media item but also a series of links from a collection of source programme components (images and video clips) to a set of destination media. More generally a *bridget* can be a collection of links from a set of source media items to a set of destination media items, e.g.
    1. An image points to an image or a set of images
    2. An object in an image points to an object in an image or to an image or to a set of images
    3. A slide part of a slide show points to audio clip
    4. An audio clip in an audio points to the corresponding score sheet
    5. Different images drawn from a programme point to different web pages
    6. A video clip from a video points to a set of related videos

A *bridget*, however, is not simply a URL, but contains two data structures: one related to the source media item and the other related to the destination media item. It may also contain information on how the *bridget* itself should be presented to the user.

Typical *bridget* workflow
Bridgets are links which exist because of some inherent semantic relationship between content items. As such, they can be products of an editorial decision, taken by someone as the result of the inspection (which can be manual or automatic) of content items, and can be objects of a workflow which involves different roles taking care of finding, organising and finally crafting the data that constitute them.

The nature of a *bridget* is however quite different than traditional linearity of media content, and as such it induces a different, more "distributed" workflow. In fact, whether a piece of media content is a candidate source or destination for a *bridget* can be the result of an editorial decision taken at any moment and by quite different kind of users.

What is foreseeable is a sort of "layered" approach at producing *bridgets* in which actors with different roles defines *bridgets* under different perspectives and possibly concurring at the same time. Authors of programmes will define *bridget* end points (i.e. sources and destination content items) following criteria matching with the editorial intention, main distribution channel or target audience of the programme. At the same time marketing and commercial operators (e.g., advertisement agents) will define such end points following their own mind-setting, which may be independent from the authorial perspective. Last, but definitely not least, final users can define their own ways for *bridgets* through social media interaction. All the above approaches can include not only the generation of the linking information but also of information related to how referenced content have to be presented graphically or should interact with the user.
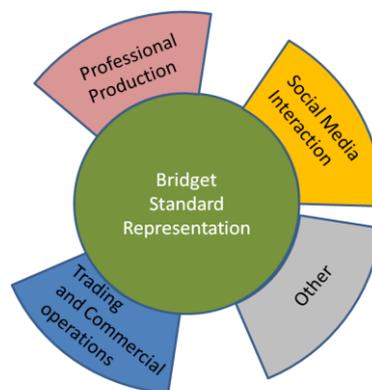
Figure 3: Bridget creation workflow.

The result is that the way in which *bridget* information would be created is quite different from traditional linear approaches to media production, and more following a daisy chain paradigm (see Figure).

Following this vision, as a natural consequence we see the prominent need to have a standard format for representing and exchanging bridget-related information in order to integrate all the systems that, in the different and disparate aforementioned domains, will have a role in generating *bridget* information.

Requirements
An updated set of requirements, that such a media linking application format implementing the *bridget* concept shall satisfy have been identified. These are subdivided in six main categories:

1. Source production data
    - Identification metadata
    - Technical metadata
2. Source content data
    - Source types
    - Access metadata
    - Descriptive metadata

3. Link related data
   - Contextual metadata (source-destination relationship)
   - Bridget production metadata
   - Segment-based linking
   - Destination content combination
   - State-varying destination content
   - User-defined flags
   - Pictorial and acoustic representations
   - Link consumption information
4. Destination content data
   - Destination types
   - Access metadata
   - Descriptive metadata
5. Presentation data
   - Scene description
   - Media
6. Identify and analyse real signals
   - Local processing
   - Remote processing

Enabling Technologies
- MPEG-21 Digital Item Declaration
- MPEG-7 Audio Visual Description Profile (AVDP)
- EBU Core Metadata Set
- MPEG-4 Augmented Reality Application Format (ARAF)

Output documents
N15096 - Requirements for Media Linking Application Format (MLAF)
N15203 - Technology under Consideration for Media Linking Application Format

3   Exploration on Media Orchestration

This exploration activity born out of the "workshop on media synchronisation for hybrid delivery" that took place in 110th MPEG meeting. As a result MPEG is exploring the use cases and requirements for new standards for the advanced synchronization of media.

With so many capture and display devices, and with applications and services moving towards a more immersive experience, we need the tools to be able to manage multiple, heterogeneous devices over multiple, heterogeneous networks, to create a single experience. We call this process Media Orchestration: orchestrating devices, media streams and resources to create such an experience.

Media orchestration:
- Applies to capture as well as consumption;
- Applies to fully offline use cases as well as network-supported use, with dynamic availability of network resources;
- Applies to real-time use as well as media created for later consumption;
- Applies to entertainment, but also communication, infotainment and professional services;
- Concerns temporal (synchronization) as well as spatial orchestration;
- Concerns situations with multiple sensors as well as multiple rendering devices, including one-to-many and many-to-one scenarios;
- Concerns situations with a single user as well as with multiple (simultaneous) users, and potentially even cases were the "user" is a machine

A couple of representative use case are given below.

Augmented broadcast (capture and rendering)
People visiting an event use their smartphones to create recordings. When the event is broadcasted, these user recordings are used to augment the "official" broadcast. For viewers watching the event at home, some of the user recordings are depicted on the main screen, as they are related directly to something shown in the main broadcast. Users at home can also use their tablet to "look around" at the event venue, where the content they see on their tablet depends on the position of their tablet. In other words, they use their tablet to browse the available sources. They can receive alerts in which area event visitors are making recordings, and can view/hear these recordings by moving their tablet around (i.e. "surround vision"). The viewers will also hear the audio that goes with the area viewed on their tablet.

Joining multiple speakers in an ad-hoc fashion to create a single speaker system
A group of users are together, and wish to play some music. They all brought their mobile devices, which they take out and place throughout the room. One user selects and start the music from a tablet, using the speakers of the various devices to create a rich sound, filling the room. The playback is controlled in such a way that the proper stereo or surround sound is reproduced, as intended by the artist. This includes adapting to specific device locations and orientation and controlling volume and equalizer settings.

Output documents
N15323 - AHG on Media Orchestration
N15091 - Context and Objectives for Media Orchestration

4   Exploration on Genome Compression and Storage

The sequencing of the genetic information of human genome has become affordable due to high-throughput sequencing technology [1], [2]. This opens new perspectives for the diagnosis and successful treatment of cancer and other genetic illnesses. However, there remain challenges, scientific as well as computational, that need to be addressed for this technology to find its way into everyday practice in healthcare and medicine. The first challenge is to cope with the flood of sequencing data. For instance, a database covering the inhabitants of a small country like Switzerland would need to store a staggering amount of data, about 2,335,740 Terabytes. The second challenge is the ability to process such a deluge of data in order to 1) increase the scientific knowledge of genome sequence information and 2) search genome databases for diagnosis and therapy purposes. High-performance compression of genomic data is required to reduce the storage size, increase transmission speed and reduce the cost of I/O bandwidth connecting the database and the processing facilities.

The current trends in sequencing data generation show clearly that the storage and transfer (bandwidth) costs will soon become comparable to the costs of sequencing. This means that IT costs may soon become a major obstacle to such genome analysis applications as personalized medicine, early diagnostics and drugs discovery, unless genetic data compression reduces IT costs on par with sequencing costs.

Document N15094 has been drafted with the goal to help MPEG to assess the opportunity to start a standardization effort in genetic information processing, particularly compression, and provides:
1. An overview of the current status of tools and technology supporting genomic information compression and storage
2. An analysis of related challenges for the stakeholders
3. A review of the existing compression tools and techniques.

Applications

Diagnostics and Personalized Medicine

Genome-based diagnostic tests have been recently developed to make personalized treatment possible thanks to discovered links between specific genetic variants and diseases. Such tests have the potential to predict risk and drive preliminary therapeutic interventions, to detect onset of disease, or detect residual disease. Although clinicians and patients are still far from being educated in how best to apply genetic knowledge in better targeting (that is, in whom to intervene) and tailoring (how best to intervene) preventive efforts, improved health is a major goal of genomic research.

Drug Discovery

Complete knowledge of the functions of all human genes might dramatically change drug discovery development processes and drug research as a whole. The application of genomic technologies to the clinical development of new and existing drugs is known as pharmacogenomics. Thanks to the recent development in genomics and pharmacogenomics in clinical research and clinical medicine, diseases could be treated in a close future according to genetic and specific individual markers, so that medications and dosages could be optimized according to the genetic profile of individual patients.

Biomarker Discovery

In medicine a biomarker is an indicator of the presence of a disease state or any other physiological state. More generally anything that is measurable and related to the state of an organism can be considered a biomarker. The pharmaceutical industry is increasingly interested in biomarker discovery because biomarkers could represent early signals of disease in clinical trials, and possible drug targets.

Genomic information generation and manipulation

Figure 4 shows the main stages of genomic information manipulation in existing bioinformatics applications. The steps depicted include:

1. Sequencing: expression of genomic information as strings (a.k.a. sequences or reads) of nucleotides identifiers.
2. Alignment/mapping: sequences arrangement to identify regions of similarity among them (*de-novo* assembly) or with respect to an external reference (a pre-constructed genome). Sequences are encoded in the form of SAM files and its binary dual named BAM [3].
3. Compression: data encoding to use less bit.
4. Storage: compressed data is stored and made available via database interfaces or files.
5. Decompression/access: access to data to perform analysis.
6. Update: previously sequenced genomic information might be updated by means of new alignment techniques or new sequencing (a.k.a. re-sequencing).

Figure 4: Genomic information generation and manipulation stages.

References

[1] S. D. Kahn, "On the Future of Genomic Data," *Science,* vol. 331, pp. 728-729, 2011.

[2] S. Wandelt, M. Bux and U. Leser, "Trends in Genome Compression," *Journal of Current Bioinformatics,* 2013.

[3] H. Li, "SAM/BAM and related specifications," [Online]. Available: http://samtools.github.io/hts-specs/.

Output documents
N15092 - Database for Evaluation of Genome Compression and Storage
N15093 - Requirements on Genome Compression and Storage (Draft)
N15094 - Investigation on Genome Compression and Storage
N15098 - Tutorial on Genome Compression
N15322 - AHG on Requirements on Genome Compression and Storage

5   Extending MVCO for Media Segments and Multi-track Rights Tracking

The Media Value Chain Ontology (MVCO), part of ISO/IEC 21000, is an ontology for formalising the representation of the Media Value Chain. This preliminary contribution (M36112) by the Centre for Digital Music, Queen Mary University of London proposes possible extensions to MVCO ontology for use cases involving intellectual property (IP) rights for temporal segments of IP entities, e.g. a song played at a given time in a podcast, as well as IP entities composed of multiple parallel tracks, such as multi-track audio. Figure 5 illustrates individual tracks of a multi-track audio resource, and a segment on its timeline.
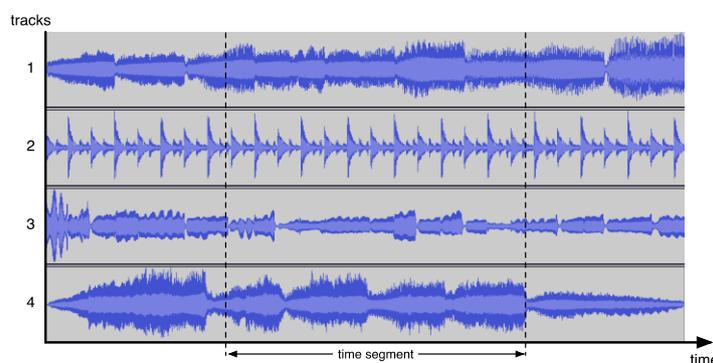

Figure 5: Visualised multi-track audio, showing individual tracks (vertical axis) and
a time segment (horizontal axis).

Use case: Podcast

A podcast, a program of music or talk made available in digital format, may be defined as a single IP entity. However, in many cases a podcast consists of a number of media items, such as music recordings with individual property rights. In order to describe these items with MVCO, the ontology needs to allow for temporal descriptions, such as "this song is played at this particular time in the podcast". The Timeline Ontology is designed to express such temporal information as depicted in Figure 6.
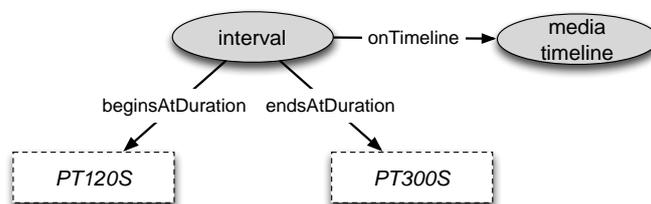

Figure 6: Describing an interval on the media timeline with the Timeline Ontology.

Use case: Multi-track audio

To enable rights tracking of multi-track audio, for instance music published in the IM AF format, additional classes and properties are defined as depicted in Figure 7.
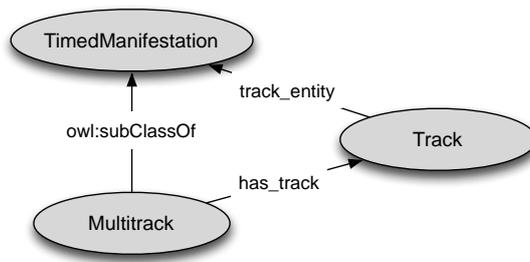
8

Figure 7: A multi-track media entity as a sublass of mvco:Manifestation

A detailed contribution towards extending MVCO with respect to the aforementioned use cases is expected for the next MPEG meeting.